

Uncovering the Risk–Return Relation in the Stock Market

HUI GUO and ROBERT F. WHITE*^{*}

ABSTRACT

There is ongoing debate about the apparent weak or negative relation between risk (conditional variance) and expected returns in the aggregate stock market. We develop and estimate an empirical model based on the intertemporal capital asset pricing model (ICAPM) that separately identifies the two components of expected returns, namely, the risk component and the component due to the desire to hedge changes in investment opportunities. The estimated coefficient of relative risk aversion is positive, statistically significant, and reasonable in magnitude. However, expected returns are driven primarily by the hedge component. The omission of this component is partly responsible for the existing contradictory results.

THE RETURN ON THE MARKET PORTFOLIO plays a central role in the capital asset pricing model (CAPM), the financial theory widely used by both academics and practitioners. However, the intertemporal properties of stock market returns are not yet fully understood.¹ In particular, there is ongoing debate in the literature about the relation between stock market risk and return, and the extent to which stock market volatility moves stock prices. This paper provides new evidence on the risk–return relation by estimating a variant of Merton’s (1973) intertemporal capital asset pricing model (ICAPM).

In his seminal paper, Merton (1973) shows that the conditional excess market return, $E_{t-1}r_{M,t} - r_{f,t-1}$, is a linear function of its conditional variance, $\sigma_{M,t-1}^2$, (the risk component) and its covariance with investment opportunities, $\sigma_{MF,t-1}$, (the hedge component); that is,

*Guo is at the Research Department, Federal Reserve Bank of St. Louis, and Whitelaw is at the Stern School of Business, New York University and the NBER. This paper was formerly circulated under the title “Risk and Return: Some New Evidence.” We thank N.R. Prabhala for providing the implied volatility data and an anonymous referee, the editor, Rick Green, and seminar participants at the Federal Reserve Bank of New York, the University of Illinois, Urbana-Champaign, the University of Texas, Austin, Vanderbilt University, the Australian Graduate School of Management, and the University of Queensland for helpful comments. Financial support from the C. V. Starr Center for Applied Economics is gratefully acknowledged by the first author. The views expressed in this paper are those of the authors and do not necessarily reflect the official positions of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

¹ The expected stock market return was long considered to be constant until relatively recent work documenting the predictability of market returns (e.g., Fama and French (1989)). It is now well understood that time-varying expected returns are consistent with rational expectations. See Campbell and Cochrane (1999) and Guo (2004) for recent examples of this literature.

$$E_{t-1}r_{M,t} - r_{f,t-1} = \left[\frac{-J_{WW}W}{J_W} \right] \sigma_{M,t-1}^2 + \left[\frac{-J_{WF}}{J_W} \right] \sigma_{MF,t-1}, \quad (1)$$

where $J(W(t), F(t), t)$ is the indirect utility function of the representative agent with subscripts denoting partial derivatives, $W(t)$ is wealth, and $F(t)$ is a vector of state variables that describe investment opportunities.² The term $\frac{-J_{WW}W}{J_W}$ is relative risk aversion, which is usually assumed to be constant over time. If people are risk averse, then this quantity should be positive.

Merton (1980) argues that, under certain conditions, the hedge component is negligible and the conditional excess market return is proportional to its conditional variance.³ Since Merton's work, this specification has been subject to dozens of empirical investigations that use widely differing specifications and estimation techniques. While these papers draw conflicting conclusions on the sign of the coefficient of relative risk aversion, in general, most studies find a weak or negative relation. Examples include French, Schwert, and Stambaugh (1987), Campbell (1987), Glosten, Jagannathan, and Runkle (1993), Whitelaw (1994), and more recent papers, including Goyal and Santa-Clara (2003), Lettau and Ludvigson (2003), and Brandt and Kang (2004). Notable exceptions are concurrent papers by Bali and Peng (2004) and Ghysels, Santa-Clara, and Valkanov (2005) that document a positive and significant relation: Bali and Peng (2004) use intraday data to estimate the conditional variance and study the risk-return tradeoff at a daily frequency, in contrast to much of the literature that uses lower frequency returns; Ghysels et al. (2005) use functions of long lags of squared daily returns to proxy for the monthly conditional variance. However, neither paper focuses on the hedge component and its interaction with the risk component as we do in this study.

The failure of the literature to reach a definitive conclusion on the risk-return relation can be attributed to two factors. First, neither the conditional expected return nor the conditional variance is directly observable; certain restrictions must be imposed to identify these two variables. Instrumental variable models and autoregressive conditional heteroskedasticity (ARCH) models are the two most commonly used identification methods. In general, empirical results are sensitive to the restrictions imposed by these models. For example, Campbell (1987) finds that the results depend on the choice of instrumental variables. Specifically, the nominal risk-free rate is negatively related to the expected return and positively related to the variance, and "these two results together give a perverse negative relationship between the conditional mean and variance for common stock" (Campbell (1987), p. 391). In the context of ARCH models, if one changes the conditional distribution of the return shock from normal to

² Strictly speaking, equation (1) is the discrete time version of Merton's ICAPM (see Long (1974)). In addition, the equation holds for the aggregate wealth portfolio for which we use the market portfolio as a proxy.

³ Specifically, if the optimal consumption function of the representative agent is much less sensitive to the state variables that describe investment opportunities than to wealth, or, if the variance of the change in wealth is much larger than the variance of the change in the state variables, then this approximation will be reasonable (see footnote 12 in his paper).

Student's-*t*, the positive relation found by French et al. (1987) disappears (see Baillie and DeGennaro (1990)).

Second, there are no theoretical restrictions on the sign of the correlation between risk and return. Backus and Gregory (1993) show that in a Lucas exchange economy, the correlation can be positive or negative depending on the time-series properties of the pricing kernel. This result suggests that the hedge component can be a significant pricing factor and can have an important effect on the risk-return relation. In general, the risk-return relation can be time varying as observed by Whitelaw (1994). The theory, however, still requires a positive partial relation between stock market risk and return. The more relevant empirical issue is to disentangle the risk component from the hedge component.

Scruggs (1998) presents some initial results on the decomposition of the expected excess market return into risk and hedge components. Assuming that the long-term government bond return represents investment opportunities, he estimates equation (1) using a bivariate exponential generalized autoregressive conditional heteroskedasticity (GARCH) model and finds that the coefficient of relative risk aversion is positive and statistically significant. However, the approach he uses has some weaknesses. For example, he assumes that the conditional correlation between stock returns and bond returns is constant; Ibbotson Associates (1997) provide evidence that it actually changes sign over time in historical data. After relaxing this assumption, Scruggs and Glabadanidis (2003) fail to replicate the earlier results. Of course, this latter result does not imply a rejection of equation (1); rather, it challenges the assumption that bond returns are perfectly correlated with investment opportunities.

Here, we develop an empirical specification based on Merton's (1973) ICAPM and Campbell and Shiller's (1988) log-linearization method and implement estimation using instrumental variables.⁴ Instead of working with the *ex ante* excess return, which is not directly observable, we decompose the *ex post* excess return into five components, namely, the risk and the hedge components, which together make up expected returns, revisions in these two components, which measure unexpected returns due to shocks to expected returns, and a residual component that reflects unexpected returns due to revisions in cash flow and interest rate forecasts. We explicitly model the volatility feedback effect,⁵ and we also control for innovations in the hedge component. Therefore, we explain part of the unexpected return on a contemporaneous basis and improve the efficiency of the estimation and the identification of the risk and hedge components of expected returns.

Another innovation relative to previous work is that we use monthly volatility implied by S&P 100 index option prices as an instrumental variable for the

⁴ French et al. (1987) argue that full-information maximum likelihood estimators such as GARCH are generally more sensitive to model misspecification than instrumental variable estimators.

⁵ Pindyck (1984, 1988), Poterba and Summers (1986), French et al. (1987), Campbell and Hentschel (1992), and Wu (2001) all emphasize the importance of the volatility feedback effect in detecting the risk-return relation.

conditional market variance.⁶ Implied volatility is a powerful predictor of future volatility, subsuming the information content of other predictors in some cases (see, for example, Christensen and Prabhala (1998) and Fleming (1998)). Implied volatility is therefore an efficient instrumental variable and improves the precision of the estimation.

We obtain three important and interesting results from estimating the model with the implied volatility data. First, the coefficient of relative risk aversion is positive and precisely estimated at 4.93 with a standard error of 2.14 in our fully specified model. Second, we find that expected returns are primarily driven by changes in investment opportunities, not by changes in stock market volatility. The two together explain 6.8% of the total variation in stock market returns, while the latter alone explains less than 1% of the variation. Moreover, other than for two short episodes associated with severe market declines, the variance of the estimated hedge component is larger than that of the risk component. Third, the risk and hedge components are negatively correlated. Thus, the omitted variables problem caused by estimating equation (1) without the hedge component can cause a severe downward bias in the estimate of relative risk aversion.

One concern about our approach is that the implied volatility data only start in November 1983. In order to check the robustness of our results, we also estimate the model with longer samples of monthly and quarterly data, in which the conditional market variance is estimated with lagged financial variables. The results from this empirical exercise can also be readily compared to those in the existing literature. Similar results are obtained in this longer data set. For the monthly data, the point estimate of the coefficient of relative risk aversion is 2.05 with a standard error of 2.98. In spite of the longer sample, the standard error is higher due to the imprecision associated with estimating the conditional variance using financial variables rather than implied volatility. For the quarterly data, the estimate of relative risk aversion is 7.75 with a standard error of 2.79. In both cases, expected returns are driven primarily by the changes in investment opportunities.

These analyses allow us to explain the counterintuitive and contradictory evidence in the current literature. The primary issue is a classical omitted variables problem. Because the omitted variable, the hedge component, is large and negatively correlated with the included variable, the risk component, the coefficient is severely downward biased and can even be driven negative. In addition, the conditional variance is often measured poorly, thus generating large standard errors and parameter estimates that can vary substantially across specifications. Finally, controlling for the effect of shocks to expected returns on unexpected returns (i.e., the volatility feedback effect and the analogous effect of innovations in the hedge component) can increase the efficiency of our estimation, sometimes substantially. However, the results suggest that care must be taken when including these effects because misspecification and estimation

⁶ Part of the implied volatility data are constructed by Christensen and Prabhala (1998) and are kindly provided to us by N. Prabhala.

error can cause the inclusion of these components to degrade the performance of the model in some cases.

The remainder of the paper is organized as follows. Section I presents a log-linear model of stock returns that decomposes ex post returns. The data are discussed in Section II. The empirical investigation is conducted in Section III, and Section IV concludes the paper.

I. Theory

A. A Log-Linear Asset Pricing Model

We first derive an asset pricing model based on Merton's ICAPM and Campbell and Shiller's (1988) log-linearization method. The log-linear approximation provides both tractability and accuracy.

As in Campbell and Shiller (1988), the continuously compounded market return $r_{M,t+1}$ is defined as

$$r_{M,t+1} = \log(P_{M,t+1} + D_{M,t+1}) - \log(P_{M,t}), \quad (2)$$

where $P_{M,t+1}$ is the price at the end of period $t + 1$ and $D_{M,t+1}$ is the dividend paid out during period $t + 1$. Throughout this paper, we use upper case to denote the level and lower case to denote the log. In addition, the subscript M is suppressed for notational convenience.

Using a first-order Taylor expansion around the steady state of the log dividend price ratio $\bar{d} - \bar{p}$, equation (2) can be rewritten as a first-order difference equation for the stock price,

$$r_{t+1} \approx k + \rho p_{t+1} - p_t + (1 - \rho)d_{t+1}, \quad (3)$$

where

$$\rho = \frac{1}{1 + \exp(\bar{d} - \bar{p})}, \quad (4)$$

$$k = -\log(\rho) - (1 - \rho)\log\left(\frac{1}{\rho} - 1\right), \quad (5)$$

and ρ is set equal to 0.997 (0.98) for monthly (quarterly) data as in Campbell, Lo, and MacKinlay (1997, Chapter 7). Henceforth, for simplicity we replace the approximation sign in equation (3) with an equals sign. Although in general the approximation error may not be negligible, Campbell et al. (1997, Chapter 7) provide evidence that it is very small in our context.

Solving equation (3) forward and imposing the appropriate transversality condition, we obtain an accounting identity for the price that also holds ex ante. Substituting this equation back into (3), we decompose the ex post stock return into two parts—the expected return and the shocks to the return (see, for example, Campbell et al. (1997, Chapter 7)). For the excess market return, $e_{t+1} \equiv r_{t+1} - r_{f,t}$, where $r_{f,t}$ is the nominal risk-free rate, this decomposition can be rewritten as

$$\begin{aligned}
e_{t+1} = & E_t e_{t+1} - \left[E_{t+1} \sum_{j=1}^{\infty} \rho^j e_{t+1+j} - E_t \sum_{j=1}^{\infty} \rho^j e_{t+1+j} \right] \\
& - \left[E_{t+1} \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j} - E_t \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j} \right] \\
& + \left[E_{t+1} \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - E_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} \right], \quad (6)
\end{aligned}$$

where Δd_{t+1+j} is dividend growth. Unexpected excess returns are made up of three components, namely, revisions in future expected excess returns, revisions in the risk-free rate forecasts, and revisions in cash flow forecasts.

Merton's ICAPM (equation (1)) provides the model for expected excess returns

$$E_t e_{t+1} = \gamma \sigma_t^2 + \lambda \sigma_{MF,t}, \quad (7)$$

where $\frac{-J_{WW}W}{J_W} = \gamma$ (the coefficient of relative risk aversion) and $\frac{-J_{WF}}{J_W} = \lambda$, which are both assumed to be constant over time. Substituting equation (7) into equation (6) and noting that $E_t e_{t+1+j} = E_t[E_{t+j}e_{t+1+j}]$ by iterated expectations, we get

$$e_{t+1} = \gamma \sigma_t^2 + \lambda \sigma_{MF,t} - \eta_{\sigma,t+1} - \eta_{F,t+1} - \eta_{f,t+1} + \eta_{d,t+1}, \quad (8)$$

where

$$\eta_{\sigma,t+1} = E_{t+1} \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2 - E_t \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2, \quad (9)$$

$$\eta_{F,t+1} = E_{t+1} \sum_{j=1}^{\infty} \rho^j \lambda_{t+j} \sigma_{MF,t+j} - E_t \sum_{j=1}^{\infty} \rho^j \lambda_{t+j} \sigma_{MF,t+j}, \quad (10)$$

$$\eta_{f,t+1} = E_{t+1} \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j} - E_t \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j}, \quad (11)$$

and

$$\eta_{d,t+1} = E_{t+1} \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - E_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}. \quad (12)$$

The first two terms in equation (8) capture the expected excess return. The third and fourth terms explicitly write out the unexpected return due to shocks to the risk component and hedge component of expected returns, respectively. The remaining terms are shocks to risk-free rate and cash flow forecasts.

B. Modeling the Risk and Hedge Components of Returns

The empirical implementation of equation (8) requires further specification of the risk and hedge components of returns. By imposing a specific time-series

model on these components, we can also reduce the shocks to these components, which are written in equations (9) and (10) as infinite sums, to more manageable closed-form terms.

First consider the risk component of expected returns and the shock to this component. To construct an empirical model of the conditional variance, we project the realized variance, v_{t+1}^2 , onto a vector of state variables, Z_t ; that is,

$$v_{t+1}^2 = \omega_0 + \omega_1 Z_t + \zeta_{t+1}. \quad (13)$$

The fitted value from the estimation is used as a proxy for the conditional market variance;⁷ that is,

$$\hat{\sigma}_t^2 = \omega_0 + \omega_1 Z_t. \quad (14)$$

For the longer sample period we use one lag of realized volatility in addition to a set of financial predictor variables as the state variables. For periods for which we have implied volatility data, we also add this variable to Z_t . Discussion of the computation of the realized variance and the choice of state variables is postponed until Section II.

In order to calculate the innovation in the risk component and its effect on unexpected returns we need to compute the shock to this conditional variance, which, in turn, requires specifying a process for the state variables. Following Campbell and Shiller (1988), among others, we assume that the state variables, Z_{t+1} , follow a vector autoregressive (VAR) process with a single lag⁸

$$Z_{t+1} = B_0 + B_1 Z_t + \varepsilon_{Z,t+1}. \quad (15)$$

Because Z_t includes the realized variance, the VAR process in equation (15) subsumes equation (13). Given this law of motion,

$$\eta_{\sigma,t+1} = \rho\gamma\omega_1(I - \rho B_1)^{-1}\varepsilon_{Z,t+1}, \quad (16)$$

where I is an identity matrix with the same dimension as the vector Z_t (see the Appendix for details). Note that the unexpected return due to revisions in the risk component is a linear function of the shocks to the state variables that define the conditional variance. This term generates the volatility feedback effect in equation (8), that is, returns are negatively related to contemporaneous innovations in the conditional variance.

There are several ways to estimate the hedge component ($\lambda_{\sigma_{MF,t}}$) in equation (8). Scruggs (1998) uses a bivariate exponential GARCH model, in which he assumes that the long-term government bond return is perfectly correlated with investment opportunities. Following Campbell (1996), we model the hedge component as a linear function of a vector of state variables X_t , that is,

⁷ This type of specification has a long history in the literature. For example, French et al. (1987) use a time-series model of realized variance to model the conditional variance. Numerous papers since then employ pre-determined financial variables as additional predictors.

⁸ Extending the VAR to longer lags is conceptually straightforward, but it adds nothing to the intuition from the model.

$$\widehat{\lambda\sigma}_{MF,t} = \phi_0 + \phi_1 X_t. \quad (17)$$

This formulation needs some explanation because, in the stock return predictability literature, it is used to model total expected returns, not just the component of expected returns due to hedging demands. The danger is that we may mistakenly attribute part of the risk component to the hedge component, that is, we will not be able to identify the two components separately. We avoid this problem by ensuring that our proxy for conditional volatility subsumes most of the information about the risk component that is contained in the state variables in equation (17). In particular, we specify the hedge component state variables, X_t , as a subset of the conditional variance state variables, Z_t . Therefore, the process of projecting realized variance on these variables guarantees that we extract all the (linear) information about future volatility that they contain. The residual predictive power that these variables have for expected returns should be due only to the hedge component.

One advantage of equation (17) is that it allows us to calculate the revision term for the hedge component, $\eta_{F,t+1}$ in equation (8), directly as in Campbell and Shiller (1988), Campbell (1991), and Campbell and Ammer (1993). By controlling for this component of returns, we can potentially increase the efficiency of the estimation and the precision with which we estimate the coefficients. Specifically, again assuming a VAR(1) for the state variables,⁹

$$X_{t+1} = A_0 + A_1 X_t + \varepsilon_{X,t+1}, \quad (18)$$

the revision to the hedge component is

$$\eta_{F,t+1} = \rho \phi_1 (I - \rho A_1)^{-1} \varepsilon_{X,t+1}, \quad (19)$$

where I is an identity matrix with the same dimension as the vector X_t . This result is analogous to the result in equation (16), with the minor exception that the shock to the risk component includes the coefficient of relative risk aversion, γ , as an additional multiplicative factor. As for the risk component, innovations to the hedge component are a linear function of the shocks to the state variables.

After substituting equations (14), (16), (17), and (19) into equation (8), we obtain the model that is estimated in this paper

$$\begin{aligned} e_{t+1} = & \gamma[\omega_0 + \omega_1 Z_t] + [\phi_0 + \phi_1 X_t] - \rho\gamma\omega_1(I - \rho B_1)^{-1} \varepsilon_{Z,t+1} \\ & - \rho\phi_1(I - \rho A_1)^{-1} \varepsilon_{X,t+1} - \eta_{f,t+1} + \eta_{d,t+1}. \end{aligned} \quad (20)$$

This equation captures the six components of excess market returns, in particular, expected returns due to the risk and hedge components, unexpected returns due to shocks to these components of expected returns, and shocks to cash flow

⁹ Because X_t is a subset of Z_t , equation (18) is inconsistent with equation (15) unless the additional variables do not enter equation (18). In implementing this procedure we guarantee consistency by zeroing out the statistically insignificant coefficients throughout. Consequently, equation (18) is subsumed by equation (15).

and risk-free rate forecasts. The risk and risk revision terms are linear functions of the estimated lagged conditional variance and the contemporaneous shocks to the state variables that define this conditional variance, respectively. The hedge and hedge revision terms are written in terms of the lagged state variables and the shocks to these variables, respectively. The shocks to the cash flow and risk-free rate forecasts are not written out explicitly, and they form the regression residual in the specification that we estimate.

II. Data Description

The model is estimated with three sets of data. The first data set utilizes the volatility implied by S&P 100 index (OEX) option prices as an instrument for the conditional market variance and covers the period November 1983 to May 2001. The implied volatility series only starts in 1983, so we also use two other data sets over longer sample periods (July 1962 to May 2001 for monthly data and 1952Q1 to 2002Q3 for quarterly data) that adopt commonly used financial variables as instruments to estimate the conditional market variance.

The implied volatility series is a combination of the data constructed by Christensen and Prabhala (1998) and the VIX data calculated by the Chicago Board options Exchange (CBOE). Christensen and Prabhala (1998) compute nonoverlapping monthly implied volatility data for the S&P 100 index spanning the period from November 1983 to May 1995. It is important to note that the S&P 100 index option contract expires on the third Saturday of each month. Christensen and Prabhala compute implied volatility based on a contract that expires in 24 days. The sampling month is thus different from the calendar month; moreover, some trading days are not included in any contract. For example, the implied volatility for October 1987 is calculated using the option price on September 23, 1987 for the option that expires on October 17, 1987. For November 1987, it is based on the option price on October 28, 1987 for the option that expires on November 21, 1987. Thus, trading days between October 17, 1987, and October 28, 1987, including the October 19, 1987 stock market crash, are not included in any contract. We extend this series to May 2001 by using the VIX, which is a calendar-month implied volatility series constructed from the options with expiration dates that straddle the relevant month end.¹⁰

One alternative to splicing these two series is to use the VIX series from its inception in January 1986. We do not pursue this alternative for two reasons. First, the Christensen and Prabhala series starts in November 1983, thereby providing an extra 2 years of data. Second, and more important, the Christensen and Prabhala series is likely to have less measurement error. Christensen and Prabhala use a single option maturity for each observation, and we match our realized variance to this maturity date. In contrast, in the VIX series each data point is actually an average of implied volatilities from options with less than a month to maturity and options with more than a month to maturity. This maturity mismatch induces measurement errors. Consistent with this intuition,

¹⁰ See the CBOE web site (<http://www.cboe.com>) for details and the historical data series.

rerunning the analysis using just the VIX series produces similar results except that the predictive power of implied volatility for realized variance is lower and the resulting estimates of the coefficient of relative risk aversion are less precise. Apart from the measurement error problem, there appear to be no significant differences between the two series; therefore, we use the merged series in the results that we report.¹¹

The monthly excess market return and realized variance are constructed from daily excess market returns. We use daily value-weighted market returns (VWRET) from the Center for Research in Security Prices (CRSP) as daily market returns. The daily risk-free rate data are not directly available. Following Nelson (1991) and others, we assume that the risk-free rate is constant within each month and calculate the daily risk-free rate by dividing the monthly short-term government bill rate (from Ibbotson Associates (1997) or CRSP) by the number of trading days in the month. The daily excess market return is the difference between the daily risk-free rate and the daily market return.

The realized monthly market variance is defined as¹²

$$v_t^2 = \sum_{k=1}^{\tau_t} e_{t,k}^2 + 2 \sum_{k=1}^{\tau_t-1} e_{t,k} e_{t,k+1}, \quad (21)$$

where τ_t is the number of days to expiration of the option contract in month t , when we are using the Christensen and Prabhala (1998) implied volatility data, or the number of days in the calendar month otherwise; $e_{t,k}$ is the daily excess market return. Equation (21) assumes a mean daily return of zero, however, adjusting the variance for the realized mean daily return over the month has no appreciable effect on the results. Equation (21) also adjusts for the first-order autocorrelation in daily returns induced by nonsynchronous trading in the stocks in the index (as in French et al. (1987)).¹³ The monthly excess market return is the sum of daily excess market returns, and quarterly returns and realized market variances are defined analogously.

For the longer sample period we estimate the conditional market variance by projecting realized variance on its own lag and two predetermined variables: (1) the consumption-wealth ratio (*CAY*) (see Lettau and Ludvigson (2001)), and (2) the stochastically detrended risk-free rate (*RREL*). The latter variable is defined as

$$RREL_t = r_{f,t} - \frac{1}{12} \sum_{k=1}^{12} r_{f,t-k}. \quad (22)$$

¹¹ Restricting the sample to the November 1983 to May 1995 subperiod spanned by the Christensen and Prabhala series also produces qualitatively similar results.

¹² See Anderson et al. (2003) and the references therein for a discussion of the theory of quadratic variation and the conditions under which realized variance is an unbiased and efficient predictor of return volatility.

¹³ In our sample, the first-order autocorrelation of daily returns is 13.9% and all higher-order autocorrelations are essentially zero.

The risk-free rate is taken from Ibbotson Associates (1997) or CRSP, and the consumption-wealth ratio is computed and supplied by Martin Lettau.¹⁴ For both sample periods we use the same two state variables, the consumption-wealth ratio and the detrended risk-free rate, to estimate the hedge component of returns.¹⁵

It is worth noting that the cointegrating vector used in computing the consumption-wealth ratio is estimated over the full sample. This methodology has been questioned, particularly in the context of out-of-sample predictability. However, our focus is on understanding the economics of the in-sample risk-return tradeoff, and we can see no apparent reason why the use of the full sample cointegrating vector will spuriously affect the estimation of this relation. The reason to choose the full sample estimate is that it greatly reduces the estimation error. Moreover, we obtain similar results with a less parsimonious specification that has more instrumental variables.

Finally, we employ three additional variables as instruments in order to calculate overidentifying restrictions tests for various models. The natural choice is a set of variables that have been shown to predict returns and/or volatility, and we use the default spread (i.e., the yield spread between Baa-rated and Aaa-rated bonds), the dividend yield, and the term spread (i.e., the yield spread between long-term and short-term Treasury securities).

III. Empirical Results

A. Econometric Strategy

We simultaneously estimate equations (15) and (20) using the generalized method of moments (GMM):

$$Z_{t+1} = B_0 + B_1 Z_t + \varepsilon_{Z,t+1}, \quad (23)$$

$$\begin{aligned} e_{t+1} = & \gamma[\omega_0 + \omega_1 Z_t] + [\phi_0 + \phi_1 X_t] - \rho\gamma\omega_1(I - \rho B_1)^{-1}\varepsilon_{Z,t+1} \\ & - \rho\phi_1(I - \rho A_1)^{-1}\varepsilon_{X,t+1} + \epsilon_{t+1}. \end{aligned} \quad (24)$$

Recall that equation (23) subsumes both equations (13) and (18) because both the realized variance and the hedge component state variables are included in the vector Z_t . Thus, the parameters ω_0 and ω_1 are rows of B_0 and B_1 , respectively, A_1 is a submatrix of B_1 , and $\varepsilon_{X,t+1}$ is a subvector of $\varepsilon_{Z,t+1}$. Throughout we first estimate an unrestricted VAR and then re-estimate the model zeroing out the statistically insignificant coefficients from the first stage. This procedure has

¹⁴ See his web site, <http://pages.stern.nyu.edu/~mlettau/>, for details.

¹⁵ An earlier version of the paper uses a somewhat different set of four instrumental variables for both the variance and the hedge component of returns: (1) *RREL*, (2) the yield spread between Baa-rated and Aaa-rated bonds, (3) the yield spread between 6-month commercial paper and 3-month Treasury bills, and (4) the dividend yield. The results are qualitatively similar, and the new specification is more parsimonious.

no meaningful effect on the estimation of the model in equation (24), but it has the distinct advantage of highlighting the key interactions among the variables.

To estimate equation (23), we use the standard ordinary least squares (OLS) moment conditions. The only subtlety is in formulating the moment conditions for equation (24). Note that theory does not imply that the terms $\varepsilon_{Z,t+1}$ and $\varepsilon_{X,t+1}$ are orthogonal to the contemporaneous regression error ϵ_{t+1} ; therefore, these variables should not be used as instruments. Using a constant, X_t , and the fitted conditional variance ($\omega_0 + \omega_1 Z_t$) as instruments is sufficient to identify the free parameters. The two shocks are functions of the fitted residuals from equation (23), the estimated parameters B_1, A_1, ω_0 , and ω_1 also come from this equation, and ρ is set to 0.997 (see Section I). The parameters γ, ϕ_0 , and ϕ_1 are identified by equation (24). What, then, is the value of including the two terms that represent shocks to the two components of expected returns? First, by reducing the amount of unexplained variation, they should improve the efficiency of the estimation and the accuracy with which the parameters of interest can be estimated. Second, these terms also depend on the parameters, and thus imposing the theoretical restrictions implied by the model may also help to pin down these parameters.

In order to understand what drives our results relative to the existing literature and to understand the gains from imposing the additional restrictions on the revision terms, we also estimate restricted versions of the model that exclude various terms in equation (24). Specifically, we consider the following five models:

$$1: \quad e_{t+1} = \phi_0 + \gamma[\omega_0 + \omega_1 Z_{t-1}] + \epsilon_{t+1}, \quad (25)$$

$$2: \quad e_{t+1} = \phi_0 + \gamma[\omega_0 + \omega_1 Z_{t-1}] - \rho\gamma\omega_1(I - \rho B_1)^{-1}\varepsilon_{Z,t+1} + \epsilon_{t+1}, \quad (26)$$

$$3: \quad e_{t+1} = \phi_0 + \phi_1 X_t + \epsilon_{t+1}, \quad (27)$$

$$4: \quad e_{t+1} = \gamma[\omega_0 + \omega_1 Z_{t-1}] + [\phi_0 + \phi_1 X_t] + \epsilon_{t+1}, \quad (28)$$

$$5: \quad e_{t+1} = \gamma[\omega_0 + \omega_1 Z_{t-1}] + [\phi_0 + \phi_1 X_t] - \rho\gamma\omega_1(I - \rho B_1)^{-1}\varepsilon_{Z,t+1} + \epsilon_{t+1}. \quad (29)$$

We denote the full model in equation (24) Model 6. In each case, we use the set of instruments corresponding to the independent variables in GMM estimation, that is, a constant, plus X_t and/or the fitted conditional variance. For example, in Model 1 there are two parameters to estimate (ϕ_0 and γ) and we use two instruments (the conditional variance and a constant). Model 2 is similar, except that the inclusion of the revision term should help identify γ . In Model 3 we use a constant and the vector of variables X_t as instruments. The remaining models use the full set of instruments. Thus, all the models are exactly identified. We also re-estimate Models 1–3 using the full set of instruments and test the resulting overidentifying restrictions. For example, the overidentified version of Model 1 also uses the vector of variables X_t as instruments. Finally, we re-estimate Models 4–6 using three additional state variables as instruments (the default spread, the term spread, and the dividend yield) and compute the resulting overidentifying restrictions test.

Table I
The Efficiency of Implied Volatility

We estimate equation (30) (a regression of realized variance, v_t^2 , on its own lag, lagged financial variables, CAY_{t-1} and $RREL_{t-1}$, and lagged implied variance, V_{t-1}^2) and restricted versions thereof over the period from January 1984 to May 2001. Heteroskedasticity-consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5%, and 10% levels are marked by ***, **, and *, respectively.

	Const.	CAY_{t-1}	$RREL_{t-1}$	v_{t-1}^2	V_{t-1}^2	R^2
v_t^2	0.020*** (0.006)	-0.043*** (0.013)	0.246 (0.186)	0.394*** (0.091)		0.274
v_t^2	0.011* (0.006)	-0.023* (0.014)	0.045 (0.163)	0.105 (0.122)	0.420*** (0.101)	0.386
v_t^2	0.057*** (0.019)				0.530*** (0.073)	0.365

^aThe constant has been scaled by a factor of 100 for presentation purposes.

B. Estimation with Implied Volatility Data

It is well known that we can predict stock market volatility with variables such as the nominal risk-free rate, the consumption-wealth ratio, and lagged realized variance (see, for example, Campbell (1987), French et al. (1987), and Lettau and Ludvigson (2003)).¹⁶ To test the information content of implied volatility, we regress the realized variance on the implied variance from S&P 100 options, V_t^2 , and these additional variables; that is,

$$v_t^2 = a_0 + a_1 X_{t-1} + a_2 v_{t-1}^2 + a_3 V_{t-1}^2 + \zeta_t. \quad (30)$$

We estimate equation (30) and restricted versions thereof with GMM. The parameter estimates and heteroskedasticity-consistent standard errors are reported in Table I.

We first exclude implied volatility in order to verify the predictive power of the other variables in our sample period. Both the consumption-wealth ratio and lagged variance enter significantly, with the signs of the coefficients consistent with the existing literature. Moreover, the explanatory power is substantial, with an R^2 of 27%. The risk-free rate does not have any marginal explanatory power, but this may be specific to our sample period. When the implied variance is added to the specification, it is highly significant and the R^2 increases to 39%. The consumption-wealth ratio remains significant at the 10% level, but the magnitude of the coefficient is reduced by approximately a factor of two.¹⁷ Finally, we also report estimates from a regression of realized variance on the

¹⁶ Other variables such as the commercial paper-Treasury spread and longer lags of realized variance have a small amount of marginal predictive power over and above CAY and $RREL$ in our sample. However, the results are invariant to including these additional variables and we restrict the set of variables for parsimony.

¹⁷ Note that throughout the paper significance levels are based on two-sided tests. To the extent that one-sided tests are appropriate in some cases, the relevant significance level is half of that reported.

implied variance alone. The R^2 declines slightly to 37%, but it is clear that implied variance is the best single predictor and that little is lost by excluding the other explanatory variables. Consequently, we select the implied variance as the single explanatory variable in the variance equation. Of equal importance, these results imply that we will be able to separately identify the two components of expected returns. The explanatory variables used for the hedge component (i.e., the consumption-wealth ratio and the risk-free rate) should pick up little of the risk component because they have limited marginal explanatory power for future variance, after controlling for the predictive power of implied volatility. Including the additional variables in the model for conditional variance has no meaningful effect on the later estimation of the full model; therefore, for ease of exposition we ignore them.

If implied variance is a conditionally unbiased predictor of future variance, then in Table I the intercept in the last regression should be equal to zero and the coefficient on implied variance should be equal to one. However, an extensive literature documents positive intercepts and slopes less than unity in similar regressions (see Poon and Granger (2003) for a survey of this literature). This bias may be related to the market price of volatility risk (see, for example, Bollerslev and Zhou (2004)). In addition, if the S&P 100 differs in an economically significant way from the value-weighted CRSP index, then the parameter estimates may also differ from zero and one. This is almost certainly the case because the realized variance of the S&P 100 index is larger than the realized variance of the CRSP value-weighted portfolio, most likely because the S&P 100 is not a well-diversified portfolio. Table I shows that while the estimated coefficient is positive, it is significantly less than one, and the intercept is significantly positive, although it is small. Thus, while implied volatility may be informationally efficient relative to other variables it is not conditionally unbiased. As a result, we use the fitted value from this estimation as our proxy for conditional variance in the estimation of the full model.

Table II reports results from the estimation of equations (23)–(24) using monthly implied volatility data for the January 1983 to May 2001 period. The results for the conditional variance process, estimated using the implied volatility data, are in the first line of Panel A, which is just the last line of Table I. The estimated processes for the state variables for both the risk and hedge components are shown in regressions 2 through 4 in the same panel. For implied volatility, the AR(1) coefficient is positive and significant at the 1% level, as expected, although the estimated degree of persistence is not very large. Both the remaining state variables, the consumption-wealth ratio and the relative T-bill rate, are quite persistent. Over this sample none of the state variables show statistically significant predictive power for their counterparts, so we zero out these coefficients for ease of presentation and interpretation. The results that follow are not sensitive to this choice.

The results from the estimation of the model for returns are reported in Panel B. Recall that we estimate six different specifications, five restricted models given in equations (25)–(29) and the full specification given in equation (24). In addition, we estimate both an exactly identified and an overidentified

Table II
Estimation Results with Implied Volatility

In Panel A we estimate equation (23) (a VAR(1) on realized variance, v_{t+1}^2 , implied variance, V_{t+1}^2 , and the financial state variables, CAY_{t+1} and $RREL_{t+1}$) with statistically insignificant coefficients zeroed out. In Panel B we estimate the model in equation (24) (a decomposition of excess returns into the risk component, with associated coefficient of relative risk aversion γ , and the hedge component, proxied by the financial state variables CAY and $RREL$) and various restricted versions thereof. The equations are estimated jointly via GMM over the period from January 1984 to May 2001. Heteroskedasticity-consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5%, and 10% levels are marked by ***, **, and *, respectively. The last two columns provide the degrees of freedom, test statistic, and p -value (in brackets) for the overidentifying restrictions test. See Section III.A for details.

Panel A: The Conditional Variance Process and the Processes for the State Variables						
	Const.	v_t	V_t^2	CAY_t	$RREL_t$	R^2
v_{t+1}^2	0.057*** (0.019)		0.530*** (0.073)			0.365
V_{t+1}^2	0.001*** (0.000)		0.641*** (0.118)			0.411
CAY_{t+1}	0.063*** (0.016)			0.859*** (0.035)		0.738
$RREL_{t+1}$	-0.002 ^a (0.003)				0.846*** (0.045)	0.702
Panel B: The Model						
	Const.	γ	CAY	$RREL$	R^2 DF	χ^2
Model 1	0.003 (0.004)	2.485 (1.909)			0.008 0	
Model 1	0.004 (0.004)	3.842** (1.918)			0.008 2	9.411 [0.009]
Model 2	0.003 (0.004)	2.485 (1.909)			0.203 0	
Model 2	0.001 (0.003)	4.571*** (1.705)			0.203 2	4.974 [0.083]
Model 3	-0.211** (0.097)		0.493** (0.215)	-4.706 (3.500)	0.034 0	
Model 3	-0.221** (0.097)		0.513** (0.214)	-4.395 (3.505)	0.034 1	4.778 [0.029]
Model 4	-0.334*** (0.103)	5.623*** (1.885)	0.744*** (0.266)	-5.767 (3.503)	0.068 0	
Model 4	-0.327*** (0.101)	5.919*** (1.731)	0.727*** (0.222)	-5.616* (3.382)	0.068 3	0.298 [0.960]
Model 5	-0.191** (0.084)	4.350*** (1.623)	0.426** (0.188)	-5.161 (3.426)	0.230 0	
Model 5	-0.201** (0.083)	3.656*** (1.352)	0.453** (0.184)	-4.279 (3.221)	0.227 3	1.038 [0.792]
Model 6	-0.189** (0.085)	4.311 (2.692)	0.424** (0.193)	-2.875 (3.794)	0.248 0	
Model 6	-0.144** (0.060)	4.928** (2.142)	0.322** (0.133)	-3.462 (3.066)	0.278 3	0.831 [0.842]

^aThe constant has been scaled by a factor of 100 for presentation purposes.

specification for each model. Model 1 is the standard risk-return model estimated in much of the literature, that is, a regression of returns on a measure of the conditional variance. However, in contrast to many existing results, we find a coefficient that is positive, albeit statistically insignificant, and reasonable in magnitude.¹⁸ If the hedge component is unimportant or orthogonal to the risk component, the coefficient value of 2.5 represents an estimate of the coefficient of relative risk aversion of the representative agent; however, this estimate may be biased downwards slightly due to measurement error in the conditional variance. The absence of a hedge component also implies that the constant in the regression should be zero—a hypothesis that cannot be rejected at the 10% significance level. However, the R^2 of the regression of less than 1% is very small, and adding the consumption-wealth ratio and risk-free rate as instruments yields a convincing rejection of the overidentifying restrictions. Even though the model can be rejected with these additional instruments, the estimate for relative risk aversion is larger and significant at the 5% level.

Model 2 attempts to refine the specification by controlling for the effect of shocks to the risk component on unexpected returns, that is, the volatility feedback effect. Adding this term leaves the parameter estimates unchanged in the exactly identified specification because it is orthogonal to the estimated conditional variance by construction. Nevertheless, the estimation does provide some corroborating evidence for the existence of a risk component in that the R^2 increases to 20%. It is not necessarily surprising that the R^2 from Model 2 greatly exceeds that from Model 1. We know from extensive empirical investigations that expected returns are a small component of returns, and therefore the explanatory power of Model 1 (and Models 3 and 4 to come) is sure to be relatively small. In contrast, Model 2 (and Models 5 and 6 to come) exploits the correlation between innovations in the state variables that describe expected returns and unexpected returns. Nevertheless, the model does put structure on the way that this correlation is exploited. Specifically, a fixed function of the innovations (based on parameters estimated in a separate set of equations) is added to the right-hand side of the regression. If the correlation is of the wrong sign or the shocks are of the wrong magnitude, adding this term can reduce the R^2 . In any case, these particular results should be interpreted with caution because the overidentified model is still rejected at the 10% significance level.

These initial results suggest two conclusions. First, sample period issues aside, improving the quality of the proxy for conditional variance, in our case using implied volatility, seems to help recover the theoretically justified positive risk-return relation. In concurrent work, Bollerslev and Zhou (2004) find a similar result using implied volatility. Ghysels et al. (2005) also present evidence consistent with this conclusion using functions of daily squared returns data to form a better proxy for conditional variance. They obtain a positive and

¹⁸ The standard errors are computed via GMM and are asymptotic. Given the sample size, there is some question as to whether asymptotic standard errors are appropriate. This question is impossible to answer definitively, but we also compute small-sample standard errors using a block bootstrap methodology. The resulting standard errors on the state variables are actually slightly lower than those reported in Table II, but the standard errors on γ are somewhat higher.

significant coefficient when regressing monthly returns on this measure over a longer sample period. Second, while there is some evidence of a positive risk-return relation, statistical power is weak, and the model can be rejected. Thus, controlling for the hedge component of expected returns may be important.

Model 3 estimates the standard return predictability regression from the literature using our two state variables. In our case, however, we interpret this regression as an estimation of the hedge component of expected returns without controlling for the risk component. The signs of the coefficients, positive on the consumption-wealth ratio and negative on the relative T-bill rate, are consistent with the results in the literature, as is the R^2 of just over 3%. The consumption-wealth ratio is significant at the 5% level, and the predictive power is substantially greater than that for the risk component in Model 1. However, using the estimated conditional variance as an additional instrument leads to a rejection of the model at the 5% level.

Under the ICAPM, both Models 1 and 3 are misspecified because theoretically both the risk and hedge components should enter the model for expected returns. Model 4 combines these two terms, and the results are positive. First and foremost, estimated risk aversion is now 5.6, significant at the 1% level. The risk-return relation is highly statistically significant and of a reasonable magnitude—a reversal of the weak and/or negative results in the literature. One natural explanation is that Model 1, and more generally, similar specifications in the literature, suffer from a classical omitted variables problem, that is, they do not control for the hedge component of expected returns. The effect of an omitted variable on the estimated coefficient of the included variable depends on the covariance of this variable with the included variable. In this case, if the covariance is negative, that is, the risk and hedge components are negatively correlated, then the coefficient on conditional variance will be biased downwards when this term is included alone. Second, including the risk component also helps identify the hedge component; the coefficient on the consumption-wealth ratio is now more significant. Third, the joint explanatory power of the two components exceeds the sum of the individually explained variations from the separate regressions—the R^2 increases to 6.8%. Fourth, the results give us added confidence that we are correctly identifying the risk and hedge components. The hedge component is positively related to *CAY*, while the conditional variance is negatively related to *CAY* (see Table I). Finally, the model cannot be rejected using the dividend yield, default spread, and term spread as additional instruments.

While Model 4 is theoretically well specified, it is possible that our identification of the components of returns can be improved by controlling for the effects of shocks to expected returns on contemporaneous unexpected returns. This issue is addressed by Models 5 and 6. These models trade off efficiency and potential specification error via the inclusion of innovations in expected returns. Which model provides the best tradeoff is largely an empirical question. In Model 5, we add the shock to the risk component of expected returns. The results are not dramatically different from those of Model 4, but a couple of observations are worth making. First, including the shock to the conditional variance can

effect both the estimate of relative risk aversion and the hedge component. In this case, γ drops from 5.6 to 4.4 for the just-identified specifications, and the coefficients on both the state variables are closer to zero. Second, controlling for some of the variation of unexpected returns can increase the efficiency of the estimation. In this case, the R^2 increases from 6.8% to 23.0%, and the standard errors on the coefficients drop by between 1% and 29%. As with Model 4, Model 5 cannot be rejected using the three additional instruments.

Finally, Model 6 also includes the shock to the hedge component in the regression. The estimated coefficients for both the exactly identified and the overidentified specifications are similar to those from Model 5, and the R^2 s are marginally higher. However, the standard errors are larger, although relative risk aversion still remains significant at the 5% level for the overidentified specification. These results are somewhat disappointing because if the hedge component were persistent (see Panel A) and a significant component of expected returns (see Panel B, Model 3), then shocks to the hedge component should explain a significant fraction of unexpected returns. Moreover, the denominator of *CAY* includes financial wealth (i.e., the level of the stock market) and thus innovations in *CAY* are negatively correlated with returns almost by construction. As such, the fact that the R^2 does not increase more is a testament to the strength of the theoretical restrictions imposed on the innovation term. The most likely explanation for these results is that the monthly consumption-wealth ratio is mismeasured. The monthly series is computed from the quarterly series via interpolation. Mismeasurement will not have a large adverse effect on the estimation of the hedge component because the consumption-wealth ratio and expected returns are persistent. However, the shock to the hedge component is unpredictable by definition, and mismeasurement of the consumption-wealth ratio could easily lead to a substantial degradation in the quality of the shock to this variable. Evidence to this effect is contained in the estimation with quarterly data that is discussed later.

In order to better illustrate the omitted variables problem and to gain some economic understanding of the results, we construct the fitted risk and hedge components of expected returns using the parameter estimates from the overidentified specification of Model 6. However, given the similarities between the parameter estimates for all four specifications of Models 5 and 6, all of these models generate the same conclusions. The two series are plotted in Figure 1 along with the NBER business cycle peaks and troughs (shaded bars represent recessions, the period between the peak of the cycle and the subsequent trough). In general, the hedge component is more variable than the risk component of expected returns, although the sample variance of the latter is larger due to two spikes in implied volatility in November 1987 and September–October 1998 (following steep declines in the market). The two series are negatively correlated, with a sample correlation of -0.41 ; thus, omitting the hedge component causes the coefficient on the risk component to be biased downwards. The magnitude of this bias depends on the covariance between the hedge component and the included variable (the conditional variance) relative to the variance of the included variable times its true coefficient (the true coefficient of relative




Figure 1. The components of expected returns using implied volatility. The risk (dashed line) and hedge (solid line) components of monthly expected excess returns for the period 3/84 to 5/01, using the implied volatility data. The estimation results are reported in Table II, Panel B, Model 6, overidentified specification. Recessions are marked as shaded bars.

risk aversion). In this sample, the covariance is -4.13 while the product of the estimated value of γ (from the overidentified specification of Model 6) times the variance of the implied variance is 14.56 . The bias is not sufficient to reverse the sign of the estimated coefficient in Models 1 and 2, but it is substantial.

From an economic standpoint, the hedge component appears to exhibit some countercyclical variation (i.e., it increases over the course of recessions), but there are only two recessions in the sample, one of which is only partially within the sample period, so this interpretation is extremely casual. The risk component exhibits little or no apparent business cycle patterns, although, as noted above, variation in this series is dominated by increases in implied volatility following large market declines. Of some interest, the hedge component is negative for substantial periods of time, implying that at these times the stock market serves as a hedge against adverse shifts in investment opportunities.

C. Estimation with Financial State Variables

The results of Section III.B go a long way toward resurrecting the positive risk–return relation, but the analysis suffers from two problems: (i) the sample period is relatively short, and (ii) it relies on implied volatility data that are not available in all periods or across all markets. Consequently, we now turn to an analysis that constructs conditional variance estimates from ex post variance

computed using daily returns and conditioning variables that include lagged realized variance and our two state variables, the consumption-wealth ratio and the relative T-bill rate.

Before proceeding to the full estimation, we first examine the variance process more closely by estimating a regression of realized variance on two lags plus the two state variables

$$v_t^2 = a_0 + \sum_{i=1}^2 a_{1,i} v_{t-i}^2 + \sum_{k=1}^2 a_{2,k} X_{k,t-1} + \zeta_t. \quad (31)$$

Realized stock market variance shoots up to 0.0755 in October 1987 and returns to a more normal level soon thereafter. The crash has a confounding effect on the estimation of equation (31), which is reported in Table III. Although the first lag of the market variance is the dominant explanatory variable in the pre-crash (9/62-9/87) subsample and to a lesser extent in the post-crash (1/88-5/01) subsample (the first two regressions reported in Table III, respectively), the second lag of market variance is also equally economically and more statistically significant in the full sample (the third regression). Not surprisingly, the R^2 of this third regression is also much lower (5% vs. 30% and 28% in the subsamples), and the sum of the coefficients on the lagged variance terms is also much lower (0.26 vs. 0.61 and 0.48). Basically, including the crash significantly degrades the predictive power of the regression over all the other months because this one observation dominates the sample in an OLS context. To reduce the impact of the October 1987 market crash, we somewhat arbitrarily set the realized stock market variance of October 1987 to 0.0190 basis points,

Table III
Variance Predictability Regressions

We estimate equation (31) (a regression of realized variance, v_t^2 , on its own lags and lagged financial variables, CAY_{t-1} and $RREL_{t-1}$) for various subsamples. Heteroskedasticity-consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5%, and 10% levels are marked by ***, **, and *, respectively.

Subperiod		Const.	CAY_{t-1}	$RREL_{t-1}$	v_{t-1}^2	v_{t-2}^2	R^2
9/62-9/87	v_t^2	0.004 (0.003)	-0.009 (0.006)	-0.009 (0.087)	0.441*** (0.117)	0.172** (0.067)	0.298
1/88-5/01	v_t^2	0.015** (0.007)	-0.032** (0.015)	0.141 (0.185)	0.268*** (0.094)	0.213** (0.098)	0.278
9/62-5/01	v_t^2	0.012*** (0.004)	-0.023** (0.009)	0.016 (0.093)	0.151* (0.077)	0.113** (0.056)	0.048
9/62-5/01 ^a	v_t^2	0.007*** (0.002)	-0.015** (0.006)	0.015 (0.078)	0.360*** (0.066)	0.207*** (0.047)	0.252
9/62-5/01 ^a	v_t^2	0.008*** (0.003)	-0.015*** (0.006)	-0.014 (0.084)	0.453*** (0.067)		0.219

^aThe market variance of October 1987 is adjusted as discussed in Section III.C.

the second largest realization in our sample.¹⁹ The corresponding results are shown in the fourth regression in Table III. The coefficients and explanatory power look similar across both the subperiods and the full sample after this adjustment, and the coefficient on the second lag of realized variance is relatively small but statistically significant (at the 1% level) in the full sample. However, this second lag does not add much to the explanatory power, as demonstrated by the fifth regression, which excludes this term. The R^2 drops only 3%, from 25% to 22%. Of some interest, the consumption-wealth ratio is a significant predictor of future variance in all but the pre-crash regression, entering with a negative coefficient. Consequently, our final specification for the conditional variance has the consumption-wealth ratio and a single lag of the realized variance, where the October 1987 variance is adjusted as described above.

Table IV reports results for the estimation of the full system in equations (23) and (24) using monthly data over the period September 1962 to May 2001. The estimation of the variance process is reported in the first line of Panel A. Other than the substantial degree of explained variation, the key result is that the conditional variance is negatively and significantly related to the consumption-wealth ratio. Panel A also reports the estimation of the process for the two state variables. The results are comparable, although not identical, to those estimated over the shorter sample period in Table II. Again, both variables exhibit strong persistence and cross-variable predictability is limited.

Panel B reports the major results of interest, the estimates from the six models in Section III.A and also those estimated using the implied volatility data in Section III.B. Models 1 and 2 contain only the risk component (the conditional variance) plus, in the latter case, the innovation in this component of expected returns. The effects of the omitted variable, that is, the hedge component, and measurement issues in the conditional variance are clear. The coefficient on the conditional variance is predominantly negative and the standard error is large. The negative coefficient is consistent with previous studies that document a negative risk-return relation over similar sample periods (e.g., Whitelaw (1994)). The fact that the standard error is 50% or more larger than in the shorter sample period (which uses implied volatility) is a testament to the value of finding better proxies for the conditional variance. Not surprisingly, using the consumption-wealth ratio and risk-free rate as additional instruments generates rejections of the model at all conventional levels.

When both components are estimated together, as in Model 4, the coefficient on the conditional variance is still negative for the exactly identified specification, and the standard error is still high for both specifications. Moreover, the estimated risk component of expected returns is small; the R^2 of 3.6% in

¹⁹ Campbell et al. (2001) suggest this type of adjustment to the crash variance, although a similar result could be achieved by dropping this observation from the sample. There may be other reasons to exclude the October 19, 1987 market crash from the sample. Schwert (1990) shows it is unusual in many ways, and Seyhun (1990) argues that it is not explained by fundamentals. It is not predicted by the option data used in this paper.

Table IV
Estimation Results with Financial State Variables

In Panel A we estimate equation (23) (a VAR(1) on realized variance, v_{t+1}^2 , and the financial state variables, CAY_{t+1} and $RREL_{t+1}$) with statistically insignificant coefficients zeroed out. In Panel B we estimate the model in equation (24) (a decomposition of excess returns into the risk component, with associated coefficient of relative risk aversion γ , and the hedge component, proxied by the financial state variables CAY and $RREL$) and various restricted versions thereof. The equations are estimated jointly via GMM over the period from September 1962 to May 2001. Heteroskedasticity-consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5%, and 10% levels are marked by ***, **, and *, respectively. The last two columns provide the degrees of freedom, test statistic, and p -value (in brackets) for the overidentifying restrictions test. See Section III.A for details.

Panel A: The Conditional Variance Process and the Processes for the State Variables						
	Const.	v_t^2	CAY_t	$RREL_t$	R^2	
v_{t+1}^2	0.008*** (0.003)	0.454*** (0.067)	-0.015*** (0.006)		0.219	
CAY_{t+1}	0.066*** (0.011)		0.851*** (0.026)		0.726	
$RREL_{t+1}$	0.002** (0.001)		-0.004** (0.002)	0.813*** (0.056)	0.672	
Panel B: The Model						
	Const.	γ	CAY	$RREL$	R^2	DF χ^2
Model 1	0.008 (0.006)	-1.459 (3.039)			0.001	0
Model 1	0.011* (0.006)	-2.917 (3.233)			0.001	2 15.130 [0.001]
Model 2	0.008 (0.006)	-1.468 (3.066)			0.001	0
Model 2	0.004 (0.004)	1.455 (2.182)			0.030	2 14.500 [0.001]
Model 3	-0.166*** (0.058)		0.390*** (0.133)	-5.827*** (2.117)	0.036	0
Model 3	-0.166*** (0.058)		0.391*** (0.132)	-5.812*** (2.097)	0.036	1 0.003 [0.958]
Model 4	-0.164** (0.068)	-0.160 (3.051)	0.387** (0.148)	-5.832*** (2.132)	0.036	0
Model 4	-0.175*** (0.065)	0.657 (2.777)	0.408*** (0.142)	-5.637*** (2.105)	0.035	3 0.925 [0.819]
Model 5	-0.164** (0.067)	-0.161 (3.070)	0.387*** (0.147)	-5.827*** (2.115)	0.035	0
Model 5	-0.167** (0.065)	-0.084 (2.898)	0.392*** (0.143)	-5.716*** (2.110)	0.035	3 0.978 [0.807]
Model 6	-0.182** (0.070)	1.597 (3.482)	0.421*** (0.153)	-5.474** (2.326)	0.040	0
Model 6	-0.179*** (0.066)	2.053 (2.979)	0.412*** (0.145)	-2.816 (1.851)	0.067	3 4.272 [0.234]

Model 4 is no higher than that of Model 3, which excludes the risk component. In contrast, the hedge component appears to be well identified with the coefficients on both state variables significant at the 5% level. It is interesting to note that the hedge component, in contrast to the conditional variance, is positively related to the consumption-wealth ratio. Thus, we obtain the negative covariance between the hedge and risk components, which exacerbates the omitted variable problem in Models 1 and 2 and also aids in the separate identification of the two components of expected returns.

Adding shocks to the risk component of expected returns in Model 5 has almost no effect on the results. In contrast, the innovation in the hedge component, included in Model 6, causes a slight improvement in the R^2 and the imprecise estimate of the relative risk aversion is of a reasonable magnitude. Given that neither the risk nor hedge components of expected returns are identified very well due to estimation error and possible misspecification, it is not totally surprising that the shocks to these components provide little assistance.

The results in Tables II and IV are not directly comparable because they cover different sample periods. Consequently, as a robustness check, we re-estimate the models in Table IV using the sample period in Table II and report the results in Table V. The Table V results confirm those of Table IV. Of course, the processes for *CAY* and *RREL* are identical to those in Table II, as is Model 3 (without the overidentifying restrictions). Of primary interest, the estimate of γ is negative over the shorter sample period when only the risk component is included, and the sign is reversed when controlling for the hedge component. Again, the estimate of relative risk aversion is reasonable in magnitude but statistically insignificant at conventional levels. As before, the innovations in the risk and hedge components do not dramatically improve the fit of the model. In other words, the qualitative nature of the results in Table IV are preserved in the subsample covered by the implied volatility data in Table II. Thus, we have confidence that the evidence in Table II is not sample specific.

Finally, we re-estimate the six models using quarterly data over the sample period 1952Q3 to 2002Q3 and report the results in Table VI. The quarterly frequency may produce better results because the monthly *CAY* series may be subject to measurement error as mentioned above, the predictive ability for returns of variables such as the consumption-wealth ratio tends to increase at longer horizons (Lettau and Ludvigson (2001)), and longer lags of realized variance can be useful in predicting returns (Ghysels et al. (2005)).

The results at the quarterly horizon confirm those of Table IV, but there is superior identification of the coefficient of relative risk aversion. The variance process and the process for the other state variables in Panel A exhibit the same features. Models 1 and 2 show that the standard risk-return regression generates a negative coefficient, as it does at the monthly frequency. In fact, the overidentified specifications generate estimates that are negative and statistically significant. However, these models are strongly rejected using the consumption-wealth ratio and risk-free rate as additional instruments, indicating the serious misspecification problem caused by omitting the hedge component.

Table V
Estimation Results with Financial State
Variables–11/83-5/01 Subsample

In Panel A we estimate equation (23) (a VAR(1) on realized variance, v_{t+1}^2 , and the financial state variables, CAY_{t+1} and $RREL_{t+1}$) with statistically insignificant coefficients zeroed out. In Panel B we estimate the model in equation (24) (a decomposition of excess returns into the risk component, with associated coefficient of relative risk aversion γ , and the hedge component, proxied by the financial state variables CAY and $RREL$) and various restricted versions thereof. The equations are estimated jointly via GMM over the period from January 1984 to May 2001. Heteroskedasticity-consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5%, and 10% levels are marked by ***, **, and *, respectively. The last two columns provide the degrees of freedom, test statistic, and p -value (in brackets) for the overidentifying restrictions test. See Section III.A for details.

Panel A: The Conditional Variance Process and the Processes for the State Variables						
	Const.	v_t^2	CAY_t	$RREL_t$	R^2	
v_{t+1}^2	0.020*** (0.006)	0.390*** (0.090)	-0.043*** (0.013)		0.269	
CAY_{t+1}	0.063*** (0.016)		0.859*** (0.035)		0.738	
$RREL_{t+1}$	-0.002 ^a (0.003)			0.846*** (0.045)	0.702	
Panel B: The Model						
	Const.	γ	CAY	$RREL$	R^2	DF χ^2
Model 1	0.013** (0.005)	-2.029 (2.639)			0.004	0
Model 1	0.014*** (0.005)	-1.564 (2.448)			0.003	2 5.851 [0.054]
Model 2	0.013** (0.005)	-2.083 (2.730)			0.032	0
Model 2	0.019*** (0.005)	-4.984* (2.470)			0.052	2 5.008 [0.082]
Model 3	-0.211** (0.097)		0.493** (0.215)	-4.706 (3.500)	0.034	0
Model 3	-0.207** (0.097)		0.485** (0.214)	-4.941 (3.493)	0.034	1 0.900 [0.343]
Model 4	-0.303** (0.146)	3.070 (3.247)	0.687** (0.316)	-4.486 (3.516)	0.039	0
Model 4	-0.305** (0.143)	2.950 (3.272)	0.690** (0.308)	-3.706 (3.342)	0.039	3 1.047 [0.790]
Model 5	-0.308* (0.156)	3.202 (3.616)	0.697** (0.336)	-5.183 (3.773)	0.001	0
Model 5	-0.194* (0.114)	-0.566 (1.993)	0.457* (0.249)	-4.368 (3.314)	0.047	3 1.801 [0.615]
Model 6	-0.311** (0.143)	3.419 (3.390)	0.704** (0.309)	-1.428 (3.683)	0.122	0
Model 6	-0.287** (0.134)	4.116 (3.037)	0.647** (0.289)	-1.856 (3.064)	0.131	3 0.495 [0.920]

^aThe constant has been scaled by a factor of 100 for presentation purposes.

Table VI
Estimation Results with Financial State
Variables-Quarterly Data

In Panel A we estimate equation (23) (a VAR(1) on realized variance, v_{t+1}^2 , and the financial state variables, CAY_{t+1} and $RREL_{t+1}$) with statistically insignificant coefficients zeroed out. In Panel B we estimate the model in equation (24) (a decomposition of excess returns into the risk component, with associated coefficient of relative risk aversion γ , and the hedge component, proxied by the financial state variables CAY and $RREL$) and various restricted versions thereof. The equations are estimated jointly via GMM over the period from 1952Q4 to 2002Q3. Heteroskedasticity-consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5%, and 10% levels are marked by ***, **, and *, respectively. The last two columns provide the degrees of freedom, test statistic, and p -value (in brackets) for the overidentifying restrictions test. See Section III.A for details.

Panel A: The Conditional Variance Process and the Processes for the State Variables						
	Const.	v_t^2	CAY_t	$RREL_t$	R^2	
v_{t+1}^2	0.003*** (0.000)	0.423*** (0.088)	-0.095*** (0.027)		0.231	
CAY_{t+1}	-0.002 ^a (0.048)		0.833*** (0.038)		0.693	
$RREL_{t+1}$	-0.001 ^a (0.014)		-0.022** (0.009)	0.711*** (0.081)	0.527	
Panel B: The Model						
	Const.	γ	CAY	$RREL$	R^2	DF χ^2
Model 1	0.024 (0.016)	-1.379 (3.107)			0.002	0
Model 1	0.057*** (0.013)	-8.712*** (2.420)			0.004	2 16.893 [0.000]
Model 2	0.025 (0.018)	-1.547 (3.435)			0.080	0
Model 2	0.071*** (0.013)	-10.658*** (2.382)			0.184	2 10.931 [0.004]
Model 3	0.017*** (0.005)		1.945*** (0.399)	-5.599*** (1.994)	0.131	0
Model 3	0.014*** (0.005)		1.946*** (0.399)	-5.240*** (1.989)	0.131	1 3.947 [0.047]
Model 4	-0.018 (0.017)	6.273** (3.082)	2.709*** (0.573)	-5.158** (1.962)	0.154	0
Model 4	-0.014 (0.016)	6.187** (2.883)	2.446*** (0.523)	-6.210*** (1.904)	0.152	3 4.615 [0.202]
Model 5	-0.024 (0.027)	7.375 (4.880)	2.822*** (0.727)	-6.091** (2.402)	0.001	0
Model 5	-0.018 (0.019)	6.433** (3.235)	2.522*** (0.565)	-6.902** (2.177)	0.006	3 2.563 [0.464]
Model 6	-0.019 (0.019)	6.515** (3.258)	2.779*** (0.586)	-3.216 (2.107)	0.429	0
Model 6	-0.027 (0.017)	7.746** (2.795)	2.970*** (0.587)	-3.091* (1.740)	0.447	3 2.903 [0.407]

^aThe constant has been scaled by a factor of 100 for presentation purposes.

In Model 3, the hedge component shows increased explanatory power relative to the monthly frequency (with R^2 s of 13.1% for both specifications versus 3.6% in Table IV) consistent with the return predictability literature. The majority of the additional explained variation comes from an increase in the coefficient on the consumption-wealth ratio from 0.39 in Table IV to 1.94 in Table VI. However, the model can be rejected at the 5% level using the conditional variance as an additional instrument. More important, putting the risk and hedge components together in Model 4 has a sizable effect on both terms. That is, the effect of omitting either term is more dramatic than in the monthly data: γ increases to 6.3 (significant at the 5% level), the coefficient on CAY increases by almost 50%, and the R^2 s climb to over 15%. The ability of the state variables in the hedge component to help identify the risk component is further evidence that their predictive power is not spurious. Moreover, this theoretically well-specified model cannot be rejected using the three additional financial predictors as instruments.

Inclusion of the innovation to the risk component in Model 5 degrades the fit of the model, although the parameter estimates are largely unaffected and the estimate of relative risk aversion remains significant in the overidentified specification. This comparatively weak performance appears to be attributable to instability in the relation between the conditional variance and the consumption-wealth ratio over the longer sample period. Excluding this variable from the variance process in Panel A of Table VI improves the fit of Model 5, but the parameter estimates are relatively unaffected.

In contrast, adding the innovation in the hedge component in Model 6 improves the fit substantially. The R^2 s increase to 42.9% and 44.7% for the just-identified and overidentified specifications, respectively, the standard errors fall, and the model produces reasonable estimates for γ of 6.5 and 7.7, which are statistically significant at the 5% level. One important lesson is that controlling for the shocks to expected returns is only useful when these innovations are measured with reasonable accuracy.

Figure 2 plots the fitted risk and hedge components (from the overidentified specification of Model 6) over the longer sample period, with the NBER business cycle peaks and troughs marked as in Figure 1. Several observations are in order. First, the hedge component is much more variable than the risk component of expected returns. The variances differ by a factor of approximately five. Second, the correlation between the risk and hedge components is -0.55 , which generates a severe omitted variables problem when the latter is excluded. The covariance between the hedge component and the conditional variance of -4.95 is larger than the value of 4.84 for γ times the variance of the conditional variance. Thus, the omitted variables bias is sufficient to drive the estimate of relative risk aversion negative in Models 1 and 2.

From an economic standpoint, the countercyclical variation in the hedge component is much clearer than in the shorter sample period. The hedge component reaches its lowest values at business cycle peaks and its highest values at business cycle troughs. This result is not terribly surprising given the dependence of the hedge component on the consumption-wealth ratio, which is known to




Figure 2. The components of expected returns using financial state variables. The risk (dashed line) and hedge (solid line) components of quarterly expected excess returns for the period 1952Q2 to 2002Q3, using instrumental variables to estimate the conditional variance. The estimation results are reported in Table VI, Panel B, Model 6, overidentified specification. Recessions are marked as shaded bars.

be a business cycle variable. Nevertheless, our decomposition allows us to give this variation a clear interpretation as variation in the ability of the stock market to hedge shifts in investment opportunities. Specifically, stocks appear to provide this hedge at the peak of the cycle, but not at the trough, when investors require compensation for holding an asset that covaries positively with investment opportunities. The hedge component can take on a negative value for extended periods, and because it sometimes dominates the risk component the total fitted expected excess return is also negative for certain periods. From a theoretical standpoint this result is possible (e.g., see Whitelaw (2000)), but it may not be intuitively plausible. However, it is important to keep in mind that these fitted values have associated standard errors, which may be quite large given the standard errors of the coefficient estimates. Moreover, while the linear specification of the hedge component appears to be a reasonable approximation, it is most likely to break down for extreme values of the independent variables and thus be less accurate for extreme fitted expected returns.

IV. Conclusion

In this paper, we estimate a variant of Merton's (1973) ICAPM, and we find a positive relation between stock market risk and return. The estimated coefficient of relative risk aversion is reasonable in magnitude (between 1.6 and 7.7 depending on the sample period, frequency, and precise specification of the

model); therefore, the power utility function appears to describe the data fairly well. The conflicting results of previous studies are due, in large part, to the fact that they do not adequately distinguish the risk component of expected returns from the hedge component. Specifically, omitting the hedge component from the estimation causes a large downward bias in the estimate of relative risk aversion due to the negative correlation between these series.

Although stock market volatility is positively priced, in most cases it only explains a small fraction of return variation. Expected returns are driven primarily by changes in the ability of the stock market to hedge shifts in investment opportunities. Many existing economic theories can explain neither why the investment opportunity set moves so dramatically nor the macroeconomic forces behind this variation. Some recent research tries to fill this gap. For example, Campbell and Cochrane (1999) address changing investment opportunities in a habit formation model. In their model, when consumption approaches the habit level, the agent becomes extremely risk averse and demands a large expected return. Guo (2004) studies an infinite horizon heterogeneous agent model in which only one type of agent holds stocks. If there are borrowing constraints and idiosyncratic labor income shocks, shareholders require a large equity premium when their borrowing constraints are close to binding. The investment opportunities are therefore determined by shareholders' liquidity conditions.²⁰ In contrast, Whitelaw (2000) generates large changes in investment opportunities by modeling the underlying economy as a two-regime process. Because regimes are persistent, regime shifts represent large movements in investment opportunities with corresponding changes in required returns.

The focus of this paper is on understanding risk and expected returns at the market level in a time-series context; however, a significant segment of the empirical asset pricing literature focuses on the cross-section of expected returns across individual securities or portfolios. Interestingly, the importance of hedging changes in the investment opportunity set at the aggregate level is also likely to have strong implications in the cross-section. In particular, if volatility is not the primary source of priced risk at the market level, then the dynamic CAPM will not hold and market betas will not be the correct proxies for expected returns in the cross-section. Clearly, this issue warrants further investigation from both an empirical and theoretical standpoint.

Appendix: Derivation of Equation (16)

The estimated conditional market variance (equation (14)) is

$$\hat{\sigma}_t^2 = \omega_0 + \omega_1 Z_t, \quad (\text{A1})$$

where the state variables follow the process

$$Z_{t+1} = B_0 + B_1 Z_t + \varepsilon_{Z,t+1} \quad (\text{A2})$$

²⁰ Aiyagari and Gertler (1999) and Allen and Gale (1994) emphasize the liquidity effect on stock market volatility.

as in equation (15). The revision of the risk component is defined as

$$\begin{aligned}
 \eta_{\sigma,t+1} &= E_{t+1} \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2 - E_t \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2 \\
 &= \sum_{j=1}^{\infty} \rho^j \gamma (E_{t+1}[\sigma_{t+j}^2] - E_t[\sigma_{t+j}^2]) \\
 &= \sum_{j=1}^{\infty} \rho^j \gamma [E_{t+1}(\omega_0 + \omega_1 Z_{t+j}) - E_t(\omega_0 + \omega_1 Z_{t+j})] \\
 &= \sum_{j=1}^{\infty} \rho^j \gamma \omega_1 [E_{t+1} Z_{t+j} - E_t Z_{t+j}].
 \end{aligned} \tag{A3}$$

The process for the state variables implies

$$E_t Z_{t+j} = (I - B_1^j)(I - B_1)^{-1} B_0 + B_1^j Z_t. \tag{A4}$$

Therefore,

$$\begin{aligned}
 \eta_{\sigma,t+1} &= \sum_{j=1}^{\infty} \rho^j \gamma \omega_1 [(I - B_1^{j-1})(I - B_1)^{-1} B_0 + B_1^{j-1} Z_{t+1} \\
 &\quad - (I - B_1^j)(I - B_1)^{-1} B_0 - B_1^j Z_t] \\
 &= \sum_{j=1}^{\infty} \rho^j \gamma \omega_1 B_1^{j-1} [Z_{t+1} - B_1 Z_t - B_0] \\
 &= \rho \gamma \omega_1 (I - \rho B_1)^{-1} \varepsilon_{Z,t+1},
 \end{aligned} \tag{A5}$$

which is equation (16).

REFERENCES

- Aiyagari, S. Rao, and Mark Gertler, 1999, "Overreaction" of asset prices in general equilibrium, *Review of Economic Dynamics* 2, 3–35.
- Allen, Franklin, and Douglas Gale, 1994, Limited market participation and volatility of asset prices, *American Economic Review* 84, 933–955.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 579–625.
- Backus, David, and Allan Gregory, 1993, Theoretical relation between risk premiums and conditional variances, *Journal of Business and Economic Statistics* 11, 177–185.
- Baillie, Richard, and Ramon P. DeGennaro, 1990, Stock returns and volatility, *Journal of Financial and Quantitative Analysis* 25, 203–214.
- Bali, Turan, and Lin Peng, 2004, Is there a risk-return tradeoff? Evidence from high-frequency data, *Journal of Applied Econometrics* (forthcoming).
- Bollerslev, Tim, and Hao Zhou, 2004, Volatility puzzles: A simple framework for gauging return-volatility regressions, *Journal of Econometrics* (forthcoming).

- Brandt, Michael W., and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economics* 72, 217–257.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.
- Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.
- Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.
- Campbell, John Y., and John Ammer, 1993, What moves the stock and bond markets? A variance decomposition for long-term asset returns, *Journal of Finance* 48, 3–37.
- Campbell, John Y., and John H. Cochrane, 1999, Force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Ludger Hentschel, 1992, No news is good news, *Journal of Financial Economics* 31, 281–318.
- Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu, 2001, Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk, *Journal of Finance* 56, 1–43.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–227.
- Christensen, Bent J., and Nagpurnanand R. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50, 125–150.
- Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Fleming, Jeff, 1998, The quality of market volatility forecasts implied by S&P 100 index option prices, *Journal of Empirical Finance* 5, 317–345.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–30.
- Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2005, There is a risk-return tradeoff after all, *Journal of Financial Economics* 76, 509–548.
- Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, 1993, On the relation between the expected value and the variance of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.
- Goyal, Amit, and Pedro Santa-Clara, 2003, Idiosyncratic risk matters!, *Journal of Finance* 58, 975–1007.
- Guo, Hui, 2004, Limited stock market participation and asset prices in a dynamic economy, *Journal of Financial and Quantitative Analysis* 39, 495–516.
- Ibbotson Associates, 1997, *Stocks, Bonds, Bills and Inflation 1997 Yearbook* (Ibbotson Associates, Chicago, IL).
- Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815–849.
- Lettau, Martin, and Sydney Ludvigson, 2003, Measuring and modeling variation in the risk-return tradeoff, Working paper, New York University.
- Long, John B., 1974, Stock prices, inflation and the term structure of interest rates, *Journal of Financial Economics* 1, 131–170.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- Merton, Robert C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.
- Nelson, Daniel B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.
- Pindyck, Robert S., 1984, Risk, inflation, and the stock market, *American Economic Review* 74, 335–351.

- Pindyck, Robert S., 1988, Risk aversion and determinants of stock market behavior, *Review of Economics and Statistics* 70, 183–190.
- Poon, Ser-Huang, and Clive W. J. Granger, 2003, Forecasting volatility in financial markets: A review, *Journal of Economic Literature* 41, 478–539.
- Poterba, James M., and Lawrence H. Summers, 1986, The persistence of volatility and stock market fluctuations, *American Economic Review* 76, 1142–1151.
- Schwert, G. William, 1990, Stock volatility and the crash of '87, *Review of Financial Studies* 3, 77–102.
- Scruggs, John T., 1998, Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach, *Journal of Finance* 53, 575–603.
- Scruggs, John T., and Paskalis Glabadanidis, 2003, Risk premia and the dynamic covariance between stock and bond returns, *Journal of Financial and Quantitative Analysis* 38, 295–316.
- Seyhun, H. Nejat, 1990, Overreaction or fundamentals: Some lessons from insiders' response to the market crash of 1987, *Journal of Finance* 45, 1363–1388.
- Whitelaw, Robert F., 1994, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Finance* 49, 515–541.
- Whitelaw, Robert F., 2000, Stock market risk and return: An equilibrium approach, *Review of Financial Studies* 13, 521–547.
- Wu, Guojun, 2001, The determinants of asymmetric volatility, *Review of Financial Studies* 14, 837–859.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.